

Designing Visualization Capabilities for IDS Challenges

Robert F. Erbacher¹, Kim Christensen², and Amanda Sundberg³

Utah State University

ABSTRACT

This paper describes our work to develop an environment and novel visualization techniques for the visual representation, exploration, and analysis of network traffic data to ease the identification and analysis of sophisticated attacks above and beyond the ability for traditional network firewalls to detect and block. The visualization techniques are geared towards aiding analysts in filtering unwanted or unneeded data in favor of data deemed more critical and more representative of the sophisticated attacks the analysts must focus their attention on. The environment provides the needed capabilities for analyzing traditional network traffic data without additional filtering, i.e., the environment itself provides the needed capabilities.

CR Categories and Subject Descriptors: H.5.2 [Information Interfaces and Presentation]: User Interfaces – Evaluation/Methodology, I.6.9 [Visualization] – Information Visualization, Visualization Systems and Software, Visualization Techniques and Methodologies

Additional Keywords: Visualization, Intrusion Detection, Graphical User Interfaces

1 INTRODUCTION

Visualization of IDS (Intrusion Detection System) related data has been gaining increased attention recently due to the rapidly increasing number of attacks, the inability to guarantee absolute safety of computer systems, the sensitivity of information stored on computer systems, and the volume of data needing analysis. While this increased attention has garnered successes, there are still many problems still needing resolution, particularly the identification of sophisticated attacks in which the attacker:

- is attempting to avoid detection
- is appearing as an insider through the application of breached passwords at another site
- is masquerading or IP spoofing
- is applying zombie PCs breached through viruses or worms.

Thus, while visualization is proving to be the critical technology for aiding in the identification and analysis of attacks, especially novel attacks, it requires novel techniques to handle the issues intrinsic to and unique to IDS related data, including:

- Scalability. IDS related data requires the analysis of enormous volumes of data. This is particularly true with sophisticated attacks in which data over the course of

days/weeks must be analyzed and correlated.

- High Dimensionality. Basic IDS related data has high dimensionality; for example consider the number of parameters available in TCP/IP headers. Additionally, many of these parameters may themselves be considered highly dimensional; for example consider examining port numbers as a primary axis then this axis incorporates 65K dimensions. In this fashion, we can consider what IP's connect to each port, the frequency or number of accesses to that port, whether that port is known to be malicious, status of the port for each local IP (open vs. closed), local IPs using that port, etc.
- Complexity and Correlations. Elements of IDS related data can not be considered in isolation. It is the correlation of events over time that distinguishes attacks of concern.
- Temporality. IDS related data is temporal in nature and long durations of time must be considered to fully expose all attacks. This is particularly true of sophisticated attacks in which the attacker is applying a low and slow paradigm of attack.

Thus, the goal of our visualization environment is to provide the tools needed by analysts to resolve these complexities. The visualizations themselves are geared towards handling the scalability, dimensionality, and temporality of the data. This allows correlations to be presented over long segments of data without the typical extensive analysis time needed for such correlation.

We provide support for multiple data sources but through the current implementation we support raw network traffic data. While many other data sources are becoming available for aid in analysis, raw network traffic data remains one of the most widely available data sources.

2 ARCHITECTURAL OVERVIEW

Our visualization environment is designed around the architecture in figure 1. The environment is design around the goal of providing an adaptable and expandable environment to handle the needs of the analyst. Currently, the main infrastructure (center of diagram) is present as well as a single visualization and two data access protocols. The data access protocols currently implemented are based off of the simpcap environment [2] developed at AFRL (Air Force Research Lab). The first is a base simpcap environment which reads libpcap files directly. The second is a virtual simpcap format which reads a mapping file pointing to multiple libpcap files to be treated as a single file. This allows for the segmentation of very large files or the correlation of logs from multiple locales simultaneously.

The environment's expandability allows for adding additional data access paradigms and formats. This will be critical for insider threats in which host-based data (in addition to network traffic data) is critical. The myriad of data sources are coordinated by the inger manager.

¹Dept. of CS, Logan, UT 84322. Robert.Erbacher@usu.edu

²Dept. of CS, Logan, UT 84322. kinch@cc.usu.edu

³Dept. of CS, Logan, UT 84322. AmandaMB@cc.usu.edu

The display manager coordinates the instantiation and execution of the actual visualization displays, based on user selections through the user interface. It also ensures interaction is passed between displays for completely linked view support.

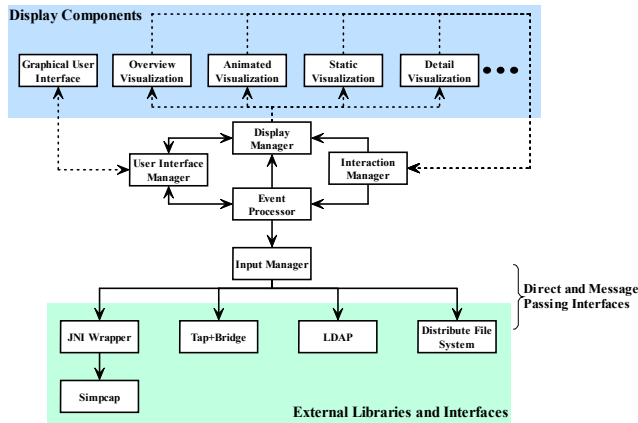


Figure 1. Architectural overview for the visualization environment.

3 USER INTERFACE

The graphical user interface (not to be confused with the direct manipulation intrinsic to the visualization displays) is designed to provide the functionality needed to allow the analyst to fully control the environment. The primary interface component is shown in figure 2. The tabs allow the user to select between different primary functional components of the interface. This initial display allows selection of the visualizations to be activated. Registering a visualization instantiates that visualization's display window. The remote control portion of this display allows the user to control the display rate. As the visualization described here is designed to be animated, the ability to control this animation is critical.

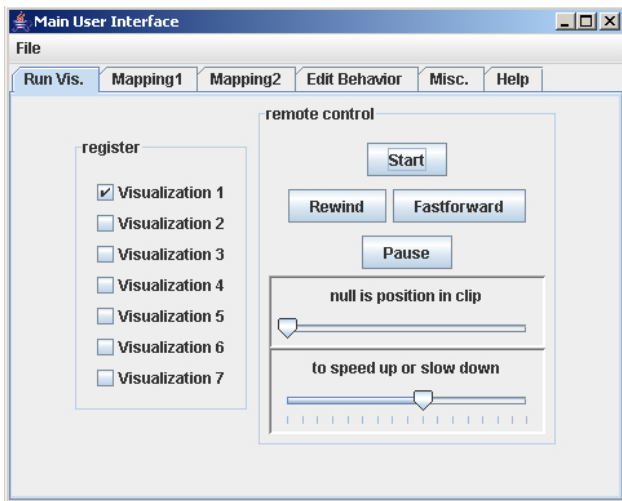


Figure 2. Primary user interface component designed to allow instantiation of visualizations and control of the execution/animation rate.

The next major component of the user interface is the data parameter to visual attribute mapping tab. This component allows the user to select which data parameters map to which visual attributes. The same data parameter may be mapped to multiple visual attributes. Such redundant mappings can aid the analysis

process. Selecting (highlighting) a data parameter and a visual attribute and clicking add, adds the selecting mapping to the systems environment.

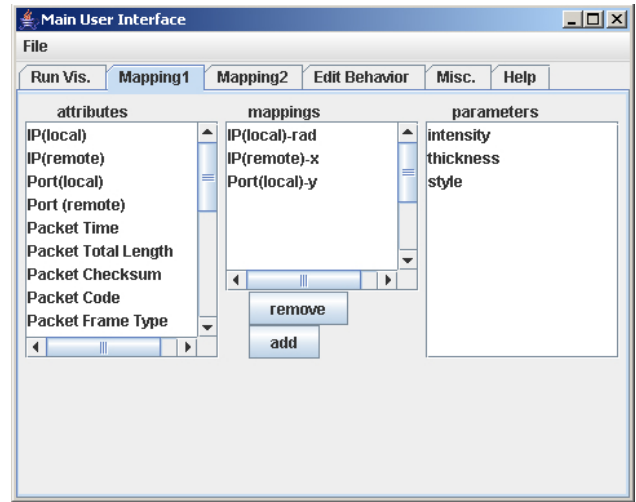


Figure 3. Parameter mapping user interface component.

Finally, the interface component in figure 4 allows the user to control miscellaneous features of the environment. This includes:

- Controlling the display of debug information.
- Setting the input file.
- Setting the gamma value. This allows the visual attributes to be brightened and is invaluable when intensity is being applied and the analyst must examine dimmer elements in detail.
- Controlling the time interval applied to the visualization elements, section 6.
- Controlling the use of GLCanvas vs. GLPanel.

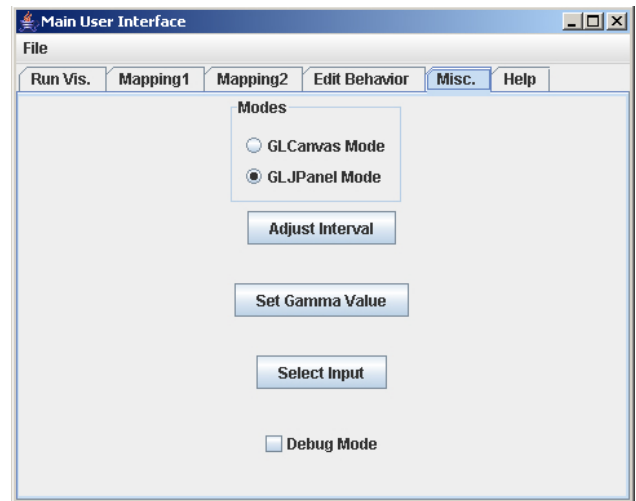


Figure 4. Miscellaneous controls for specifying environment configuration and behavior.

The last option is a necessity as this environment is developed in Java using the JOGL (Java Bindings for OpenGL) library for the Visualization techniques. JOGL currently allows developers to choose between two Windowing primitives. GLCanvas is designed to work in isolation from Swing and provides for maximum performance. GLJPanel is designed to work in a more integrated fashion with Swing but suffers a significant

performance penalty. As GLCanvas is not designed to work with Swing, an integrated environment using Swing and GLCanvas will suffer user interface anomalies, particularly in the display of the user interface components; i.e. interface components will not be updating correctly leading them to flash annoyingly.

4 THE VISUALIZATION TECHNIQUE

The visualization technique currently implemented in this environment is exemplified by the diagram in figure 5. The idea behind this technique is to represent network activity as efficiently and concisely as possible. This is to handle as much activity as possible and begin to resolve some of the many scalability issues inherent to IDS data.




Figure 5. Basic diagram of network activity monitoring visualization technique.

The developed visualization technique in its default form begins by representing the local IP address around the radius of the internal circles. The technique then represents remote IP addresses along the top and bottom of the window edges. The redundancy aids in reducing clutter and line crossings. The top edge is used if the local IP appears in the top semi-circle and the bottom edge is used if the local IP appears in the bottom semicircle. Similarly, port numbers are represented on the left and right edges of the window. If the local IP appears in the right semi-circle then the right edge of the screen is used and if the local IP appears in the left semi-circle then the left edge of the screen is used. In order to represent time periods while simultaneously maintaining persistence, we incorporate multiple rings to be representative of the age of the identified activity. The outer rings represent the most current data while each inner ring represents data m time units older, where m is the ring number. The duration of a time unit (a time interval) is user specified, see figure 4.

As indicated this is with respect to the default behavior. As discussed previously, the user is able to change the parameters mapped to each of these display features. The behavior of the visualization itself remains intrinsically the same.

An example of this visualization technique in action is shown in figure 6. This image shows an example applying actual network traffic data acquired from the Air Force Research Lab in Rome, NY. As can be seen from this display the rings associated with older data have reduced intensity to reduce their impact while maintaining the information for the analyst. Line crossings other than crossing ring boundaries are significantly limited, not only in number but also in visual impact. A future goal will be to reduce these even further.

Analysis of this display shows that there is consistent activity from several remote hosts. This display is limited in that it does

not indicate if this activity is related and the duration of time represented is still limited. The first enhancement to the technique is described in the next section and aims at identifying whether this activity is related or not.




Figure 6. Basic visualization technique showing network activity from a raw pcap file.

5 CONNECTION VS. PACKET-BASE TECHNIQUES

One of the major differences between various techniques for the analysis or visualization of IDS related data is whether said analysis applies to packet-based data or connection-based data. Connection-based data has the advantage in that it already provides a significant correlation of the data which is needed for complete and effective analysis. It also reduces the amount of data needing analysis significantly, without loss of information. It can, however, miss important information, such as port scans, and more importantly low and slow port scans. Our environment allows the user to select between packet-based and connection-based representations. As we allow analysis of raw packet data, the packet-based mode is a direct representation of the available data. The connection-based mode performs the correlation within the visualization environment to associate activity to a connection. A connection in this scenario is identified by the following characteristics:

- Identical local IPs
- Identical remote IPs
- Identical Local Ports
- Identical Remote Ports
- Packet interval < specified timeout value

Thus, in this scenario we have the advantage in that we have all the connection-based information represented as succinctly as possible. Additionally, packet-based information such as port scans will be represented as well. The mode of operation is selected in the "Edit Behavior" tab, as shown in figure 7. Figure 8 shows the same data as in figure 6 except using this connection-based metaphor. Clearly this data set contains very little connection-based information. However, this representation does validate that the heavily active remote IPs are not maintaining a single connection but many short connections. Examining the

actual animation of the visualization highlights this fact as the port numbers continuously shift upwards. The connection-based mode does reduce the amount of memory and analysis required to examine connection-based data, however, significant attention is required to identify such activity. This would be particularly true for low and slow scans.

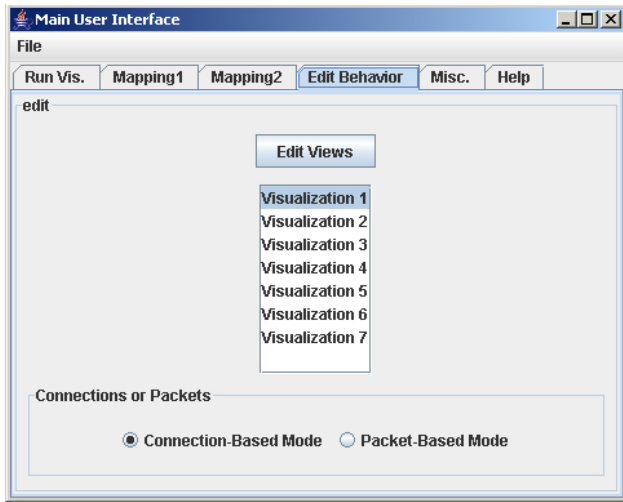


Figure 7. Edit behavior tab and associated mode selection. Allows selection between pure packet-based and connection-based representation




Figure 8. Basic visualization technique showing network activity from a raw pcap file. This representation applied the connection-based metaphor.

6 MODIFYING THE TIME SCALE

The next enhancement to the environment is geared towards handling the temporal scale of the data. In particular, the time units of the rings described previously are intrinsically linear. In other words, we can represent x time units in which x is the

number of rings. While the duration of a time unit is adjustable, making the time unit too long will inhibit analysis of current activity and temporally segregating and analyzing the data. The temporal association of this data is too critical to be cast aside so readily. Thus, we have begun to explore other scales. As exhibited in figure 9 we can change the scale from its default linear scale to an exponential one.

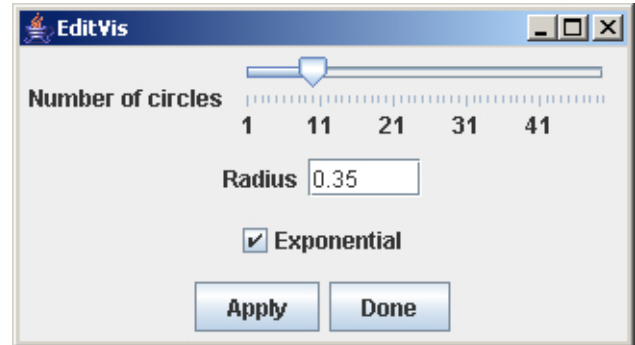


Figure 9. Edit behavior dialog. Achieved from the edit button in figure 7.

The impact of this exponential scale is exhibited in figure 10. Clearly this will allow the representation of far more data, and far longer durations of time, than would otherwise be possible. This enhancement will also have the advantage of collating information over time. It is this enhancement that will provide analysts with the ability to identify low and slow scans as the inner rings will accumulate the activity making said activity for more readily observable with far less effort than would otherwise be required.




Figure 10. Basic diagram of network activity monitoring visualization technique. This representation is modified to provide an exponential time scale for each of the rings as opposed to the default linear scale shown in figure 5.

The same data shown in figures 6 and 8 is represented again in figure 11, this time using the exponential scale. The exponential scale clearly accumulates the data. Even further, it can be clearly seen that there are three port scans occurring to multiple local hosts in rapid succession. The fact that two of the overlapping scans appear in different rings maintains the temporal information that would otherwise be lost. Examining such data using traditional techniques would take far longer. Additionally, the exponential techniques allow for the accumulation of data over very long periods. Thus, the exponential technique will be effective even for low and slow scans indicative of a sophisticated

attacker and of which an analyst or systems administrator would want to be aware of.




Figure 11. Basic visualization technique showing network activity from a raw pcap file. This representation applies the connection-based metaphor and the exponential time scale for the rings.

7 ENHANCING FIDELITY

The visualization techniques provide, in essence, temporal information through the ring-based metaphor in conjunction with either the linear or exponential scale. Switching between these modes greatly limits the range and detail to which the user can focus their attention.

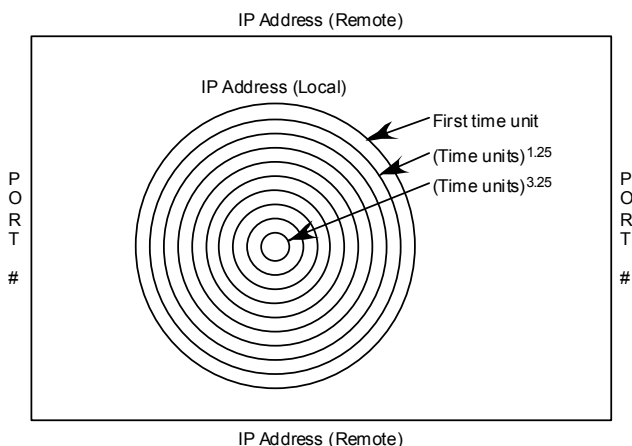


Figure 12. This representation is modified to provide an exponential time scale for each of the rings with the exponent increasing by .25 for each ring as opposed to the default linear increment shown in figure 10.

In order to improve the fidelity of the visualization we allow the user to modify two additional characteristics of the visualization. First, the exponent increment can be specified. This is the “adjust

Interval” value identified in figure 4. Rather than limiting the increment to 1 the increment can be set to any positive fractional value. Figure 12 shows an example of the implication of this change.

Second, the number of rings can be increased or decreased, figure 13. Increasing the number of rings allows the exponent increment to be decreased while still preserving the accumulation of temporal information. In addition, it allows, with either a linear or exponential scale, the user to more finely examine the temporal relationships between elements. Combined, these two techniques provide the analyst with the ability to tightly control how the display responds to temporal activity and improves the analyst’s ability to comprehend and interpret the temporally-based activity.

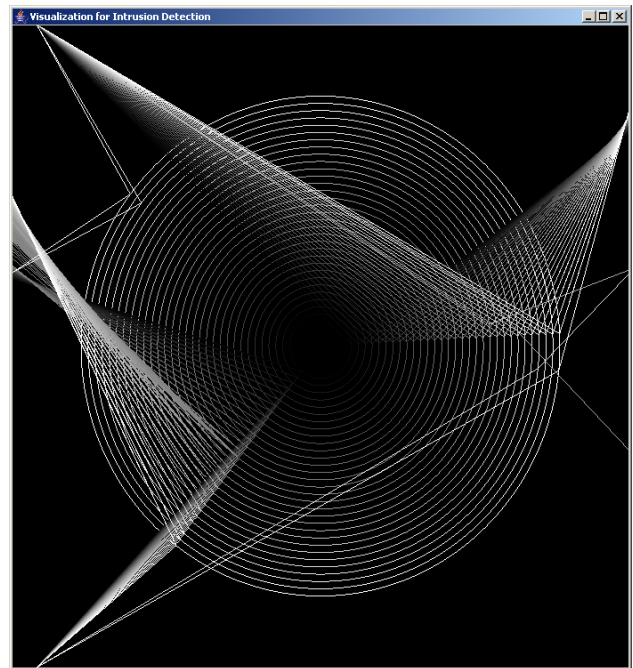


Figure 13. Basic visualization technique showing network activity from a raw pcap file. The number of rings is greatly increased allowing for greater fidelity.

8 RELATIONS TO PRIOR WORK

Currently, system administrators must analyze log files; data generated by other intrusion detection tools, and network traffic data, to analyze an attack. Other intrusion detection tools include portstentry [18][31], tripwire [10], COPS [8], BlackIce [17][27], and many others. The resultant log files can incorporate millions of messages per day. The amount of data available results in system administrators not fully collecting or monitoring all available information for systems under their control, rather focusing on primary systems and examining minimal information from the remaining systems. In fact, system administrators in general do not collect any data related to Microsoft or Apple based operating systems; even though it is possible for these systems to be the target of break-ins, subversion, and misuse. Network traffic itself is only analyzed intermittently. Ultimately, the goal must be to identify an attempted break-in or attack before the attack is successful so that the situation can be monitored and a response initiated before harm occurs. Additionally, this monitoring and analysis must be done at reduced cost, requiring fewer analysts to perform the monitoring activity.

While many intrusion detection tools have begun to incorporate basic graphical user interfaces (BlackICE [27], RealSecure [28],

Cisco Secure IDS [29], eSecure [30]) they fall short of providing effective visualization displays to aid in interpreting the generated information. For example, most of the tools will provide an indication when it received an unexpected packet. But was this an attack, a misdirected packet, a casual attack, or a real attempt to break into the system? These systems do not adequately provide the detail and event interrelationships needed to analyze the activity in the detail needed forensically.

In terms of visualization, many intrusion detection environments will incorporate “odometer-like” scales or apply other techniques to represent system state [23]. This is embodied in the Hummer “perceived level of threat” [16] indicator. Earlier systems, such as DIDS [11][19], provided graphical representations in the form of color to indicate when a system had experienced a sequence of suspicious events. While useful, these approaches do not provide adequate information to aid diagnosis.

8.1 Visualization systems

Historically, visualization has been applied extensively to network monitoring and analysis, primarily for monitoring network health and performance [1][3][9][12], even down to the router [3], individual packets [7], and individual e-mail messages [6]. The techniques developed for these purposes do not provide sufficient detail or handle sufficient numbers of nodes and attributes in combination for our needs. Initial visualization techniques for IDS environments focused on simple scales and color representations to indicate state or level of threat [4][5][16][23]. Other environments provide basic graphical user interfaces but fall short of providing the needed visual capabilities for analysis [27][28], [14][29]. The few visualization techniques that have been developed for intrusion detection have been limited as to their applicability and effectiveness.

The work by Teoh et al [20][21] focuses on Internet routing data and thus is limited in its applicability in intrusion detection and will have no applicability to forensics. The work by Eick et al. [6] strictly deals with e-mail and subsequently resolves many fewer nodes and attributes than is needed for intrusion detection. These environments typically deal with small numbers of processors that are working on a single task and thus have a common grounding and have not been applied to intrusion detection. Many environments are geared towards naïve monitoring of port activity, Teoh et al [15][21], but such work can not handle the scale of real data and infrastructures, differentiate sophisticated (e.g., low and slow) attacks, or handle the diversity of data and attack types that are truly exhibited in today’s environments.

The work by Yin et al. [25] and Lakkaraju et al. [22] focus on the representation of netflows and associated link relationships. Such techniques are critical for analyzing attacks and IDS data but these two techniques quickly suffer scalability issues and are limited as to the number of representable parameters.

9 CONCLUSIONS

The developed capabilities clearly show activity critical to an analyst in an exploratory analysis environment. In particular, we have focused the first visualization technique within this environment on the difficult problem of detecting sophisticated attacks, in this case a low and slow port scan. The novel visualization not only allows for the accumulation of temporally associated events, critical for this analysis, but also allows careful analysis of current and past events with a user controlled range of acuity.

The current environment provides for easy expandability either through the incorporation of additional visualization techniques or through the addition of additional data protocols; thus creating a

test bed for rapid experimentation of additional visualization techniques.

The current visualization technique provides benefits over other techniques such as parallel coordinates [13] by more distinctly representing and segregating each axis and placing the axes of primary analyst concern (local IPs) at the center of the user’s focus.

While there was concern about the performance of a java-based environment, the use of JOGL ensures hardware acceleration is employed (at least in the GLCanvas case). In addition Java itself has improved in performance such that we can achieve acceptable performance levels in getting the data to the graphics drivers.

10 FUTURE WORK

Clearly, additional visualization techniques must be incorporated. These additional visualization techniques will enhance the range of attacks able to be analyzed. As the environment grows a linked view paradigm will be followed to ensure effective analysis across displays.

This additional analysis will also require the incorporation of additional data paradigms. While the research group we were involved with at AFRL relied solely on network traffic data and this is the primary data source available to most administrators, additional data sources are becoming more readily available, useable, and necessary for attack analysis. This is particularly true of insider threats which require host-based data to identify and analyze in most scenarios. While there exists a large body of work aimed at correlating these disparate alert logs based upon clustering, probability, and similarity to predefined attacks [24][26] these existing techniques are insufficient for the scale of datasets, particularly in terms of size and duration, needing analysis for complete identification of all sophisticated attacks.

While the focus of the current environment is on the identification and analysis aspects of attack detection there is the need in the future for assistance in aiding the analyst in applying the response component. For example, many tools such as portsentry [18][31], will automatically block a remote IP identified as having instigated a naïve port scan. Similarly, we will need a simplified method for blocking hosts as having performed a sophisticated port scan. Performing such a task remotely requires careful consideration of security protocols that have long plagued SNMP (Simple Network Management Protocol).

11 ACKNOWLEDGEMENTS

Portions of this work were supported on AFRL’s summer faculty research program. The presented data was provided by AFRL during the term of the summer faculty research program as well.

REFERENCES

- [1] Richard Becker, Stephen Eick, and Allan Wilks. “Graphical methods to analyze network data,” In *IEEE International Conference on Communications ICC ‘93 Proceedings* Geneva, Switzerland, pp. 946-95, May 1993.
- [2] M.W. Corley, M.W. Weir, K.Nelson, and A.J.Karam, “Simplified Protocol Capture (SIMPCAP),” *Proceedings from the Fifth Annual IEEE SMC Information Assurance Workshop*, West Point, NY, pp. 176- 182, 2004.
- [3] Kenneth Cox, Stephen Eick, and Taosong He, “3 geographic network displays,” *ACM Sigmod Record*, Vol 25, No. 4, pp. 50, December 1996.
- [4] Anita D’Amico and Mark Larkin, “Methods of visualizing temporal patterns in and mission impact of compute security breaches,” In *DARPA Information Survivability Conference and Exposition (DISCEX II’01)*, Vol. 1, pp. 343–354, 2001.

- [5] Hervé Debar and Andreas Wespi, "Aggregation and correlation of intrusion-detection alerts," *In Recent Advances in Intrusion Detection*, pp. 85–103, 2001.
- [6] Stephen G. Eick and Graham J. Wills, "Navigating Large Networks with Hierarchies," *In Visualization '93 Conference Proceedings*, San Jose, California, pp. 204–210, October 1993.
- [7] Deborah Estrin, Mark Handley, John Heidermann, Steven McCanne, Ya Xu, and Haobo Yu, "Network Visualization with Nam, the VINT Network Animator," *IEEE Computer*, Vol. 33, No. 11, pp. 63–68, November 2000.
- [8] J.A. Fore, "System security: when enough is not enough," *Proceedings of Expanding Expectations in Integrated Online Library Systems (IOLS)*, New York, 1997, pp. 53–62.
- [9] Taosong He and Stephen G. Eick, "Constructing Interactive Visual Network Interfaces," *Bell Labs Technical Journal*, Vol. 3, No. 2, pp. 47–57, April–June 1998.
- [10] G.H. Kim, E.H. Spafford, "Writing, supporting, and evaluating Tripwire: a publically available security tool," *Proceedings of the 1994 USENIX UNIX Applications Development Symposium*, April 1994, pp. 89–107.
- [11] C. Ko and D. Frincke and T. Goan et al., "Analysis of an Algorithm for Distributed Recognition and Accountability," *ACM conference on Computer and Communication Security*, Vol. 1, No. 1, 1993.
- [12] Eleftherios E. Koutsofios, Stephen C. North, Russel Truscott, and Daniel A. Keim, "Visualizing Large-Scale Telecommunication Networks and Services," *Proceedings of the IEEE Visualization '97 Conference*, IEEE Computer Society Press, San Francisco, CA, pp. 457–461, 1999.
- [13] A. Inselberg, "The plane with parallel coordinates," *The Visual Computer*, Vol. 1, pp. 69–91, 1985.
- [14] Kiran Lakkaraju, Adam J. Lee, and William Yurcik, "Nvisionip: netflow visualizations of system state for security situational awareness," *In Proceedings of CCS Workshop on Visualization and Data Mining for Computer Security, ACM Conference on Computer and Communications Security*, October 29, 2004.
- [15] Jonathan McPherson, Kwan-Liu Ma, Paul Krystosek, Tony Bartoletti, Marvin Christensen, "PortVis: A Tool for Port- Based Detection of Security Events," *Proceedings of CCS Workshop on Visualization and Data Mining for Computer Security, ACM Conference on Computer and Communications Security*, October 29, 2004.
- [16] Polla, D., J. McConnell, T. Johnson, J. Marconi, D. Tobin, and D. Frincke, "A Framework for Cooperative Intrusion Detection," *21st National Information Systems Security Conference*, pp. 361–373, October 1998.
- [17] D. Rae and D. Ludlow, "Halt! Who goes there? [Internet intrusion detection benchtest]," *Network News (UK Edition)*, February 16, 2000, pp. 31–37.
- [18] Joel Scambray, Stuart McClure, and George Kurtz, *Hacking Exposed: Network Security Secrets and Solutions*, 2nd Edition, Osborne/McGraw Hill, 2000.
- [19] Snapp, S. et al., "DIDS (Distributed Intrusion Detection System) Motivation, Architecture and An Early Prototype," *National Information Systems Security Conference*, 1991.
- [20] S.T. Teoh, K.L. Ma, and S. F. Wu, "Visual exploration process for the analysis of internet routing data," *In Proceedings of the IEEE Conference on Visualization 2003*, 2003, pp. 523–530.
- [21] S.T. Teoh, K.L. Ma, S. F. Wu, and X. Zhao, "Case study: Interactive visualization for internet security," *In Proceedings of the IEEE Conference on Visualization 2002*, 2002, pp. 505–508.
- [22] Alfonso Valdes and Keith Skinner, "Probabilistic alert correlation," *In Recent Advances in Intrusion Detection*, pp. 54–68, 2001.
- [23] Vert, G., J. McConnell, and D. Frincke. "Towards a Mathematical Model for Intrusion," *21st National Information Systems Security Conference*, pp. 329–337, October 1998.
- [24] Alex Wood, "Intrusion detection: Visualizing attacks in ids data," *Giac practical, SANS Institute*, February 2003.
- [25] S. F. Wu, S.T. Teoh, K.L. Ma, and X. Zhao, "Case study: Interactive visualization for internet security," *In Proceedings of the IEEE Conference on Visualization 2002*, pp. 505–508, 2002.
- [26] Xiaoxin Yin, William Yurcik, Michael Treaster, Yifan Li, and Kiran Lakkaraju, "Visflowconnect: netflow visualizations of link relationships for security situational awareness," *In Proceedings of CCS Workshop on Visualization and Data Mining for Computer Security, ACM Conference on Computer and Communications Security*, October 29 2004.
- [27] <http://www.networkice.com/>
- [28] http://www.iss.net/securing_ebusiness/security_products/intrusion_detection/index.ph
- [29] <http://www.cisco.com/univercd/cc/td/doc/pcat/nerg.htm>
- [30] <http://www.esecurityinc.com/>
- [31] <http://www.psonic.com/products/>